FEDERAL TRADE COMMISSION

INDEX

COLLOQUY SESSION	PAGE
(LEAD BY:)	
MR. SALSBURG	4
MS. BUSH	40

FEDERAL TRADE COMMISSION

In the Matter of:)
REPORT TO CONGRESS PURSUANT TO)
CAN-SPAM ACT.) Matter No. P044405
--------------------------------)

THURSDAY
FEBRUARY 26, 2004

Room 294
Federal Trade Commission
600 Pennsylvania Ave., N.W.
Washington, D.C. 20580

The above-entitled matter came on for conference, pursuant to agreement at 1:10 p.m.

APPEARANCES:

ON BEHALF OF THE FEDERAL TRADE COMMISSION:

DANIEL SALSBURG
COLLEEN ROBBINS
SHERYL DREXLER
MICHELLE CHUA
JULIE BUSH
STEPHEN WARREN
LOUIS SILVERSIN
Federal Trade Commission
6th Street and Pennsylvania Avenue, N.W.
Washington, D.C. 20580-0000

PARTICIPANTS (VIA TELEPHONE):

JOHN LEVINE
YAKOV SHAFRANOVICH

P R O C E E D I N G S

MR. SALSBURG: We're ready to begin then. Yakov, I was mentioning that we have a court reporter here. The court reporter is transcribing the conversations that we have so we have the ability to cite to it when we're preparing our report to Congress. There are some formalities that I'll begin with first.

Today is Thursday, February 26, it's one p.m. Eastern Time. Today we're meeting with John Levine and Yakov Shafranovich.

MR. SHAFRANOVICH: Shafranovich.

MR. SALSBURG: Did I pronounce that correctly?

MR. SHAFRANOVICH: It's actually Shafranovich.

MR. SALSBURG: Shafranovich, okay. The purpose of the meeting is to discuss a possible National Do Not E-mail Registry. A little bit later on in the conversation, we may be joined by some of our other FTC colleagues who may ask questions about a possible bounty system that the CAN-SPAM Act also asked the FTC to study.

Because the meeting is being transcribed by a court reporter who doesn't have the benefit of seeing you, the first couple times that you speak, if you can just identify who you are, and I'm pretty sure she'll pick up pretty quickly which one of you is speaking after that.

MR. LEVINE: Okay.

MR. SALSBURG: John and Yakov, could you identify the names of your firms and role in the Internet Research Task Force?

MR. SHAFRANOVICH: Yakov Shafranovich. Well, my company is Solid Matrix Technologies Incorporated, and we basically are a business consulting firm. My role is one of the chairs of the Anti-spam Research Group of the Internet Research Task Force, and the purpose of the ASRG and the IRTF is to provide research and pre standard work for the Internet Standards community, mainly of the Internet Engineering Task Force.

MR. SALSBURG: John?

MR. LEVINE: Yes. My company is called Taughannock, T-A-U-G-H-A-N-N-O-C-K, Networks. It's a sole proprietorship. I write books about the Internet, and I consult the news and do software design, and I'm here in the role as the other co-chair of the ASRG.

MR. SALSBURG: Great. As you're both aware, the CAN-SPAM Act among other things, requires the FTC to prepare a report to Congress that sets forth a plan and timetable for establishing a National Do Not E-mail Registry. This report also, in addition to setting forth a plan and timetable, is supposed to include an explanation of any practical, technical, security, privacy, enforcement or other concerns that the Commission may have with such a registry.

This report is due in Congress on June 16 of 2004 which means we're quickly trying to gather as much information as possible so we can begin writing the report and have it be as thorough a report as possible. The meeting with you today is to help us with accomplishing that task.

Have either of you seen the Request for Information that the FTC issued on Friday regarding the registry?

MR. LEVINE: This is John. I have.

MR. SHAFRANOVICH: I haven't had a chance to look at it yet.

MR. SALSBURG: Okay. The Request for Information is a request to potential vendors to provide possible registry models and how they would go about setting up a registry. The RFI proposes a few such models and then invites any other creative possibilities that are out there to be submitted as well.

We thought it might be most useful to go through some of these models with you and see what your thoughts are in terms of the effectiveness, security and privacy and enforceability concerns you might have with these models.

So why don't we start with the first model, which is very similar to the Do Not Call Registry for telemarketing that the Commission operates. Under this model, a consumer would submit his or her e-mail address to the FTC. That e-mail address would be placed in a database. The database of registered e-mail addresses would be made available to e-mail marketers who would then scrub their mailing lists to remove the e-mail addresses of any consumer appearing on the list.

Do you have thoughts on such a model?

MR. LEVINE: Yeah. This is John. I don't think a single address model like that is workable, and it's for a couple of reasons. One is that I think it would be extremely difficult to keep such a list secure, even if the FTC provides a list of scrubbing services itself or it went through a small set of trusted vendors.

Spammers can triangulate. They could send in huge lists of e-mail addresses and then compare the scrubbed lists with the original list to figure out what addresses were removed. So the first issue there is the security issue.

The second is an issue of effectiveness. An important difference between e-mail addresses and phone numbers is that you can easily enumerate all the possible phone numbers in the U.S. You cannot easily enumerate all the e-mail addresses, and as a matter of fact, you can't even easily enumerate all of the e-mail addresses for a single person. Two examples of that are in my case I have an entire domain -- johnlevine.com. Every single address that is johnlevine.com is me, even addresses that have never been used before. Many companies have address servers that accept possible approximate addresses, so that if somebody's official e-mail address is john.smith@company.com, it might well also accept jsmith or j.smith or if the middle initial is Q, johnqsmith or jqsmith or any of a hundred variations.

And for Do Not E-mail Registry to be effective you would have to register all of those. I can come up with a bunch of other scenarios where there are many, many addresses corresponding to one person, so for these reasons -- these are the basic reasons that I think a registry of single addresses is unlikely to be workable. Yakov?

MR. SHAFRANOVICH: Yeah.   I would like to suggest that the amount of data you're reporting is much bigger than for the phone registry. The size of the data will be enormous, so that's something you will also want to take into account also, and like John mentioned, you want to provide apparently the ability for being able to list an entire domain or a list of names, not just single domains because there's just a lot of possibilities in the e-mail world that are not present in the regular world.

MR. SALSBURG: John, you mentioned that there were other scenarios where a person might have multiple addresses.

MR. LEVINE: Yeah.

MR. SALSBURG: Can you give me some other examples?

MR. LEVINE: Yes. A common scenario is sub addresses. Although my regular address is johnl@taush.com, any address of the form John L dash something is also me, and it turns out that sub addressing feature, it's a standard feature of a lot of mail systems, so that there are a lot of people that don't realize they have sub addresses, and again if you're going to -- sub addresses they've never used would still be their addresses so if you were going to -- if they were going to opt themselves out, they would have to opt out of every single possible sub address.

It's just impossible because there are literally billions of sub addresses possible for each individual e-mail address.

MR. SALSBURG: Do e-mail programs enable you to turn off that sub addressing system?

MR. LEVINE: They do, although it's extremely useful. It would be a big operational issue for me to do that. The way I use it, every time I provide an e-mail address to a web site or mailing address or to someone I don't know very well, I give them a unique address, and by using those individual sub addresses, I can both sort the mail that's coming in, and if someone provides it improperly to a third-party, I can figure out who leaks it.

So it's a very useful feature that, although it's possible to turn off. It would be a hardship to do so.

MR. SALSBURG: Do you have any sense of how many regular consumers use this feature?

MR. LEVINE: Well, the question isn't how many of them use it, the questions is how many of them have it available. My local ISP down the road, in fact, has sub addresses, and although almost none of its users use the sub addresses, if a marketer simply invented a sub address, it would be deliverable.

So that that would be a very easy way for them to circumvent this Do Not E-mail List by inventing deliverable addresses that the customer wouldn't have thought to opt-out.

MR. SHAFRANOVICH: One other thing I wanted to mention, people that have multiple addresses such as someone has a work and personal address, and he has permission to opt-out of the e-mail address, but the work address doesn't do it. It doesn't belong to him. It belongs to his company, and I don't know how you're going to deal with that issue.

I hope that you have a single e-mail registry for single e-mail addresses. Who has the permission to opt-out for who?

MR. SALSBURG: John, you began your description of concerns you had with the single address model as being the security issue and you mentioned triangulation.

MR. LEVINE: Yes.

MR. SALSBURG: Are there ways that a list could be kept secure?

MR. LEVINE: I think -- I've been thinking about it for awhile. I simply don't see anyway you can avoid the triangulation problem because the whole point of a Do Not E-mail Registry is to remove addresses from lists, and if spammers can present addresses at all, then they can use this triangulation attack.

You can avoid some other issues by not distributing a list in plain text and by distributing hashed versions, but the triangulation attack depends on the basic function of the list. No, it's unavoidable.

MR. SALSBURG: That's because ultimately the marketer gets a copy of something that allows them to figure out what on their list isn't on the registry?

MR. LEVINE: Yeah. I suppose you might try to come up with a scheme where the marketer doesn't even do the mailing and the trusted third-party does the mailing. I think that's impossibly cumbersome. Even so, there are ways using things like web bugs to guess fairly reliably which addresses were delivered and which weren't, and we're back to triangulating.

MR. SALSBURG: I think we're going to get to that third-party issue soon, so why don't we put that on hold for a bit.

MR. LEVINE: Sure.

MR. SALSBURG: You also mentioned hashing, and if a list were hashed, would that prevent hackers from getting into the registry?

MR. LEVINE: That makes it less -- if a list was hashed, that makes it less useful to steal the list per se since you can't usually take an individual hashed entry and reverse it. On the other hand, it doesn't do anything about the triangulation attack or in that case straightforward dictionary attack.

The spammer takes the most humongous list of e-mailers he can hack, he can find, he hashes them all, and he simply compares the hashes he came up with with the ones on the list. And the ones that match; he's now found some fraction of the people on the list.

Again it helps security some, but it doesn't address the fundamental problem.

MR. SALSBURG: How important -- I'm sorry, go ahead.

MR. SHAFRANOVICH: Hashing is a standard security issue procedure. The passwords are usually hashed, so if you have something that's been subject to an attack, your local database from being hacked, someone coming up with the data, that's the only thing its protecting.

MR. LEVINE: Yeah, it doesn't protect against triangulation and dictionary attack. It only protects against theft of individual entries, but in this case since there's so many entries, the statistical attacks will get some of the entries, which will still be very useful for spammers.

MR. SALSBURG: This is going to seem like a basic question, I'm sure, but can you explain why a spammer would bother to engage in a dictionary attack or a triangulation attack?

MR. LEVINE: They have -- I do not purport to have a unique insight into the psychology of spammers, but I've heard plenty of cases of Do Not E-mail lists -- I'm sorry, of Do Not Call lists, of industry Do Not Call Lists being stolen and used as a prospect list on the perverted theory that, Oh, they must get fewer phone calls so they would be better prospects.

I'm entirely confident that if some chunk of the FTC's list became available, that some spammers would have a theory like that, Oh, these will be live addresses, and they don't get everybody else's spam so they're good prospects for me.

MR. SALSBURG: Is there anything about the value of a list of valid e-mail addresses versus a list of valid phone numbers that would make an attack on a Do Not E-mail Registry more valuable or more likely to be engaged in by a spammer than an attack on a Do Not Call Registry?

MR. SHAFRANOVICH: Neither.

MR. LEVINE: Both have more data, and one difference is that we all know what all the possible phone numbers are. You go and look up a list of telephone prefixes and you know what all the phone numbers are, but there's no equivalent master list of all the possible e-mail addresses.

So that's a way to discover e-mail addresses that you couldn't find any other way, and there's no e-mail equivalent to sequentially dialing.

MR. SALSBURG: What are your thoughts on how effective such a list could be in terms of enforcement?

MR. LEVINE: With the limited tools that are made available by CAN-SPAM, not very. I mean the closest analogy we have is the Junk Fax Law, and although the FTC -- sorry, the FCC has done good enforcement against the very large violators, the most effective use of the TCPA has been individual suits against individual junk faxers.

And lacking some sort of remedy like that, I think it might be somewhat useful against the most egregious violators. It might be somewhat useful for sort of more or less legitimate bulk e-mailers that voluntarily wanted to keep themselves legal, but I don't think it would be terribly effective. I don't think any of these would be terribly effective without stronger remedies than we have available now.

MS. ROBBINS: What do you mean by stronger remedies?

MR. LEVINE: Than we have available now.

MS. ROBBINS: But what types of remedies are you envisioning?

MR. LEVINE: Oh, private right of action by recipients. It's not so much we need larger remedies. I'm not even considering putting them in jail for a thousand dollars. I want broader remedies so that individual recipients have the right to do something about it.

MR. SHAFRANOVICH: Yes, it's really a question of who they're able to sue. The Commission or the agencies or whoever is suing has limited amounts of funding. The more abilities for the Attorneys General to sue and people to sue, then it's more likely that a spammer that actually goes into the registry will get sued.

The other concern is that this will not be effective unless sufficient funding is provided for enforcement, and I don't know how much funding Congress has provided so far, but unless enough funding is provided in order to support this, whichever way you're enforcing it, nothing is going to happen.

MS. ROBBINS: Do you have any concerns about the enforceability of this in terms of actually identifying the spammers, as opposed to just how money is being funneled to enforcement? To clarify, technically, how to actually find the spammers and enforce the law that way?

MR. LEVINE: I don't see that as being an overwhelming problem. If you look at the spam suits that have been filed so far by AOL and Earthlink and so far, most of them start by filing against John Doe defendants, but they have -- but there's enough clues both on the spam and from where -- particularly if they have ordered some of the stuff the spammers are advertising and have figured out who cashed the check.

It's certainly pretty quick to turn to John Doe charges into actual defendants. No, I don't see that as a big problem.

MR. SALSBURG: What's the impact of the international nature of spam on the effectiveness of a registry and its enforcement?

MR. LEVINE: So long as the law is written so that the beneficiary of the spam is responsible for it, again I don't see that as a big issue. If it's a constraint, the technical community we've already established will just go offshore, but if you look at the actual spam you're receiving, even the stuff that's sent from Asia, the majority of it is clearly sent on behalf of American spammers who are American businesses. The spam is in English. They're trying to sell stuff that's of interest to Americans, and it's my understanding is that by and large, if they're selling goods, the goods are shipped from the U.S. The only significant Internet industry that I know that's moved offshore is gambling, which is sort of a special case.

MR. SHAFRANOVICH: I would also add that when you sign up for the registry, whatever law Congress has to pass to do that, who are you going to be targeting? Are you going targeting the person that actually sends the spam or the person that hired him?

If are you going after the actual person that sent the e-mail message out, that could be some third-party. If you find the person that hired them, that person is in the United States.

MR. SALSBURG: If there were to be a single address model registry, about how many registrations do you think would be made?

MR. LEVINE: Oh, man.

MR. SALSBURG: How big a database are we looking at?

MR. LEVINE: Well, if you're looking at the number -- if you expect everybody to behave themselves and just register the addresses that they actively use, you're certainly talking about hundreds of millions.

My guess is that some people who feel exasperated and have catchall domains like I do will say, Well, if they want me to register every possible address, okay, I can do that, and you may end up with semi-automated but entirely legitimate registrations of millions and millions of addresses from an individual person or for a small network, all of which are real, but none of which have been used yet.

So that could inflate it, so the total size -- the total size of the database you have will certainly be hundreds of millions and particularly if you have people registering a lot of their variant addresses just in case. It could easily be up in the billions. It would be a very large database.

MR. SHAFRANOVICH: I believe Washington State actually has some kind of registry, which you can possibly look at the numbers that they were getting and extrapolate from there as well.

MR. SALSBURG: All right.

MR. LEVINE: That's a good idea.

MR. SALSBURG: Let's move on to the second possible model that's been discussed, and that's a domain wide registry. Domains, including ISPs, could register their domains as not to receive any marketing e-mail. What are your thoughts on this model?

MR. LEVINE: Why don't you go first, Yakov.

MR. SHAFRANOVICH: Well, we saw something like that -- it's a very technical proposal that tries to do some kind of non-soliciting type of thing, when a person comes up with a name, they say do not solicit a name. That's either going to be less data, less numbers, but the problem that I see is in theory the entire domain, if it's the domain who made the decision for everybody's address, that means the individual person won't receive whatever he wants. If he wants to receive mail, he won't be able to make that choice.

MR. LEVINE: Actually I guess I would divide this into three categories. The first scenario is a model where the domain owner is simply sending in the names of the domain and you put together a giant list. That's somewhat more workable than the set of e-mail addresses because the number of domains is a lot less.

We're talking about probably tens of millions instead of hundreds of millions, and the idea that Yakov had, again which is where you actually distribute the list, where each domain owner publishes on its mail server or along with his DNS information a no soliciting tag. I think that could be pretty workable.

I'm also a member of the CAUCE, C-A-U-C-E, the

Coalition Against Unsolicited Commercial E-mail, and in 1998 we published a proposal along those lines with a sample code everybody agrees worked.

The issue that individual users in a domain couldn't un-opt themselves out I don't see as very compelling and for two reasons. One is that the recent proposals that Yakov is referring to is one that Carl Malamud submitted to the IETF, and I helped him work on it, and it actually has a varying version where you can actually write individual addresses, but I think more importantly, ISPs are not common carriers. Network operators are not common carriers, and they actually do have the right to decide for the entire network what the rules are.

If individual people want particular kinds of mail, they can always sign up for it. And if there's a demand for sending spam lists, I'm sure that a wide variety of people will be happy to provide them. It doesn't seem like -- it doesn't seem like a major issue, particularly since there are so many different ISPs, so many mail providers, that it does not seem to me to be an onerous requirement on someone if they don't like their current ISPs' mail policies, to point out that they always have the ability to get additional addresses and additional domains.

MR. SALSBURG: You mentioned that the idea of domain owners putting a no soliciting tag in their information is more workable.

MR. LEVINE: Yes.

MR. SALSBURG: You said it seemed to be working on the ones using it. I'm sorry?

MR. LEVINE: No, go ahead.

MR. SALSBURG: What has been the experience of domains that put such a tag on? Have they really gotten no spam?

MR. LEVINE: Of course not because there's no enforcement. One thing that -- CAUCE has always been dedicated towards lobbying for activating spam laws, and this basically was our version of the best way to create a registry, but a registry is of no use if there aren't sanctions for failure to use it.

The Direct Marketing Association has a Do Not Spam List which is completely useless, and my friend, Rodney Joffe, did an experiment called EMPS which that I think you're familiar, but again technically it worked fine, but nobody used it.

MR. SHAFRANOVICH: Enforcement is an issue. That's what it comes down to. It comes down to that.

MR. LEVINE: Yeah, and again your Do Not Call List is fabulously effective, but the reason it's so much more effective than the DMA's Do Not Call List is because people have to use it, and there are sanctions if they don't.

MR. SHAFRANOVICH: I would also add and go back to that model. A federated model, which is basically what you're talking about, offers a spam registry where Do Not E-mail Lists lets each domain owner specify his setting in the registry, whatever it is, is more likely to scale. The problem you had before is you come down to getting those addresses.

If you have had some kind of a system further where a company can do it. Each domain owner do it, then that is more likely to scale.

MR. LEVINE: Exactly. I think we can confidently say we know it would scale because it basically will be one extra item of data added to the DNS which already contains the delivery address for each domain so basically every domain now has some number of what are called MX records, would add one more record with their spam policy.

That's not a large addition to the DNS. I think we can be pretty confident that there would be no scaling problems, and it would also be much cheaper to administer because nobody would have to build a gigantic central database.

MR. SALSBURG: What's a scaling problem? What do you mean by that?

MR. LEVINE: Oh, it's the costs -- there are a lot of -- for pretty much any kind of technical problem, there are a lot of approaches that look like they work when you try them on a few examples, but then you say, Okay, this worked great on ten examples, now will it work on ten million examples, and the answer is no because at that much larger scale, there are issues that you don't have when you're just doing it on an experiment.

MR. SALSBURG: Can you give me a layman's explanation of how this distributed data in the DNS registry would work? If I was a marketer, would I send an initial query, or how would I go about determining whether or not, let's say, AOL was a no spam ISP?

MR. LEVINE: Oh, yeah, you would make a query, depending on how it was implemented, either to AOL's DNS service or AOL's mail server, which would then send back a piece of information that saying AOL's spam policy is so and so, send us spam or don't send us spam.

Once they have that, then they can hold on to it, and they know that that policy will apply to all the AOL addresses in their list.

MS. ROBBINS: So would there be any change in the filters then, or is this solely a marketer just complying because they want to comply?

MR. LEVINE: Oh, it would be incumbent on the senders of mail to comply with this. In fact if you -- if every time a sending -- every time a sending program contacts AOL's mail server, the mail server sends them a threatening looking legal notice which of course nobody currently reads. Let me just tell what you it says.

It says: "America On Line and its affiliated companies do not authorize the use of its proprietary computers and computer networks to accept, transmit or distribute unsolicited bulk e-mail sent from the Internet."

So they've been putting a notice like this on every single piece of mail they accept for years, but as we've been pointing out, there's no legal sanctions on mailers if they ignore it. They have been ignoring it.

MS. ROBBINS: With the CAN-SPAM Act, there's the opt-out provision. What is your sense of how marketers are complying with the CAN-SPAM Act?

MR. LEVINE: I'm trying to think if I've seen any actual CAN-SPAM compliant mail.

MR. SHAFRANOVICH: I've seen one piece.

MR. LEVINE: Yeah. Well, of the mail that actually asks for it, most of it is now compliant, and most of it has a personal mailing address.

As far as the mail I haven't asked for, yeah, I might have seen one or two pieces, but in general if the question is whether marketers -- whether spammers are complying with the Act, is no, they're not.

MS. ROBBINS: So what makes you think they would comply with this type of system?

MR. LEVINE: In the absence of more effective enforcement, they wouldn't.

MR. SHAFRANOVICH: I believe there are a bunch of other companies that are complying with the CAN-SPAM Act when you came out with federal compliance. The bottom line is enforcement. If you enforce it, whatever law you have, if it is enforced, it will work. If there's no enforcement, then it will not work.

MR. LEVINE: Yeah. Under the current circumstances, the only marketers I could see likely to use a Do Not E-mail system would be like large banks that are not sending unsolicited ads now but figure they could get away with it if they had a good list washing system like this would provide.

MR. SALSBURG: So it actually may increase the amount of spam?

MR. LEVINE: It could since it would give more of an air of legitimacy to it and it would be much easier for them to say, Gosh, if you don't want spam, tell the FTC, and we'll stop spamming you.

MR. SALSBURG: Do either of you have any other thoughts on the domain wide system?

MR. LEVINE: I mean, if you're going to implement a Do Not E-mail List at all, I think a domain -- I think the distributed domain wide system with the notice either being on the mail server and on the DNS is by far the most workable, both technically and administratively.

MS. ROBBINS: Aside from your example about the tag, if it was just a domain wide opt-out without having that tag, how do you think that kind of system would deal with permission based e-mail and transactional e-mail, if there was such a registry?

MR. LEVINE: It shouldn't affect it because it's up to the sender to know when they have to obey the tag.

MS. ROBBINS: I'm saying in the absence of a tag, if it was a domain wide opt-out where the domain's registered, their name is on the list.

MR. LEVINE: The sender presumably knows whether he's sending transactional mail or if he's sending unsolicited ads, and my assumption would be that a Do Not E-mail List would only apply to unsolicited e-mail, not the transactional mail.

MS. ROBBINS: Okay.

MR. SALSBURG: Let's move on to another possible registry model, and that would be a model involving a register of authenticated senders. There are a number of ways that could be done, but let me throw out one possible way and let me hear your thoughts.

Under this model, an e-mail marketer would register with the Commission, obtain a registration number and enter in information regarding the IP addresses and the domains from where they're going to be sending their unsolicited commercial e-mail from.

That data, the domain and IP address, would be made available to the ISP, and the e-mail marketer would have to include the registration number in the e-mail that they send.

So, in other words, the ISP would have the registration number and access to the Commission database that had the matching IP address and domain names. Do you follow that?

MR. LEVINE: I follow it. I have to say I don't see much point to it since all the ISPs I know would simply use that list of IP addresses as a list of addresses from which they will never ever accept mail, so there wouldn't be much of an incentive for a marketer to register for it.

I mean, I can see that if you believe in a world where there are people eager to get unsolicited e-mail ads, this would be a way to get them delivered better, but everybody I know doesn't want any unsolicited e-mail ads at all, so I don't see much benefit to anyone of building a system like that.

MR. SHAFRANOVICH: Can I ask, what exactly would the purpose of such a system be?

MR. SALSBURG: Let's change the facts slightly, and instead of it being required of senders of unsolicited commercial e-mail, a requirement for any commercial e-mailer, so the purpose of it would be to insure delivery of your messages if you wanted to get them through.

MR. LEVINE: There are, in fact, some private systems that do that now. Ann Mitchell's ISIPP, is working on something like that. That can be useful as a way for a legitimate mailer to prove it's bona fide, but I don't see any reason that the FTC would want to get involved with that since that is a system where it is of direct advantage to the mailer to register. Private registries can serve that function perfectly well.

MS. ROBBINS: Do you think that kind of model could help with enforcement?

MR. LEVINE: Possibly, although the kind of people who would register there would probably be ones who would behave themselves anyway, but other than simply being a way to make it easier to tell that somebody probably wasn't worth investigating, I don't see as much of a need for enforcement.

MR. SALSBURG: Let's say I didn't register and I sent along my spam without a registration number. When the ISP goes and checks the database, there's no information on me.

MR. LEVINE: Yes.

MR. SALSBURG: Is it likely that my mail is going to be filtered and never make it to an in-box?

MR. LEVINE: I frankly don't see that it would make any difference to the situation we have now, where the ISPs are diligently trying to filter all the spam now, and they would continue to do so, so I don't see this making any difference.

MR. SALSBURG: How does this -- I'm sorry, go ahead, Yakov.

MR. LEVINE: In the absence of a number, it will look like any other spam, so it's not going -- it will be just like it is now. Yakov?

MR. SHAFRANOVICH: I mean, I've been thinking about it. I'm trying to figure out. The purpose of creating a spam list, how does the database come in? I'm kind of looking at it. How would such a database come in at all?

MR. LEVINE: Presumably the idea is that all of the legitimate spammers, if there is such a thing, would register and then you can say, Ah, anyone who hasn't registered, if they sent you spam, is an illegitimate spammer, but I would say that the ability -- I think there are much more direct ways to do the same thing, and in particular, I think that registering all the marketers is a backwards way to go.

The marketers who we want to hear from identify themselves directly to the recipients, and I just don't see any advantage of trying to put the FTC in the middle of that process.

MR. SALSBURG: A technical question for you. Can the originating IP address on a piece of e-mail be forged?

MR. LEVINE: There has been a lot of argument about that. My belief is the answer is no. There is some minor -- there's some minor exceptions. It's what's known -- as I forget what it's called.

MR. SHAFRANOVICH: It's called BGP spoofing.

MR. LEVINE: Yes, there's what's called the BGP spoofing which is basically where a bad guy tells his

ISP to route a little bit of the Net to him and which is then forwarded off to the rest of the Net, and then he gets somebody else's IP addresses for awhile and then withdraws it.

I haven't see very much of that, and those sorts of things are so disruptive to the Net in general that I don't see much of that happening.

MR. SALSBURG: What was that called again?

MR. LEVINE: BGP spoofing. The only spoofing that I've actually heard of is what's called triangular routing which is unrelated to the triangulation I referred to before, which AOL has observed lately, where basically the bad guy has on both -- on the same computer he has a fast connection through which he sends out his spam, and he has a dial-up connection to AOL.

And he puts the IP address of the dial-up connection on all of the mail going out through the fast connection, so that the return packets come in through the AOL connection. This actually works, and the point of doing this is that the only addresses that people will see are the AOL connections, and when AOL knocks him off, he then takes his next stolen AOL credit card and moves to there.

AOL has been looking at this. This turns out to be a problem that they can easily fix by adjusting some of their own filtering rules a little bit, so it's a minor problem, but I don't see it as having much effect.

So I think the short answer to your question is, I can see theoretical ways that IP spoofing as possible, but I don't see it as a large scale problem.

MR. SHAFRANOVICH: Yeah. In this country it would be premature. We've heard the idea of people that will opt-out. In theory it's possible. In practice, it's highly unlikely. For all practical reasons, in the end it cannot be spoofed.

There is another thing that has happened sometimes, the IP addresses are stolen where a spammer goes to the registry and claims to be a company A and asks them to reassign an address to him, but that's not in theory being spoofed. It's basically the ownership that is being stopped.

MR. LEVINE: That's not a technical attack. That's fairly a social or a business attack.

MR. SALSBURG: Are you familiar with other authentication proposals that have been floating out there such as Microsoft's Coordinated Spam Reduction Initiative or Yahoo!'s Domain Keys and AOL's SPF?

MR. LEVINE: Yes, we've been talking to all of them.

MS. ROBBINS: Do you have any thoughts on the efficacy of any of those models?

MR. LEVINE: They all show promise. The SPF in particular and some proposals relative to that have actually been forwarded through the ASRG, and there will be an informal session in Seoul, Korea, which I guess is coming up in two weeks that's going to comment on them.

We're all attempting to do the same thing which is to make it easier to determine that a piece of mail is actually coming from the place that it purports to be coming from.

So it will deter some kinds of forgery. It will make phishing, that's phishing spelled P-H-I-S-H-I-N-G, to try to steal people's account information a little harder, but it's not directly useful against spam since if the spammer puts a true return address that he controls on the spam, then it will pass all those tests, and it will be -- as far as they're concerned it will be legitimate.

MR. SALSBURG: In your experience, do spammers usually put true return addresses?

MR. LEVINE: No. No, because there's been no advantage for them to do so. On the other hand, these days the majority of spam is sent through hijacked machines typically on consumer DSL or consumer cable modem, and they can, with no trouble at all, put a return address corresponding with the network where the hijacked machine is and to feed basically any of these schemes.

MR. SHAFRANOVICH: The basic premise, as I mentioned before, the IP address is currently one thing that cannot easily be bought. These proposals add an extra layer to it by trying to make sure that the domain name from which the mail comes from, also cannot be spoofed.

That's the entire purpose. It's a way to add additional information. Whether it's effective or not effective is not -- we don't know. I'm just wondering something, why the Federal Trade Commission -- what's the connection to these proposals? These proposals are more of a standards method or private method. I'm just trying to figure out why they would be interested.

MR. SALSBURG: Well, for one thing, we're trying to get as much of an understanding of the current status of spam and anti-spam technology as possible, so that whatever proposal we give to Congress can be enlightened.

MR. SHAFRANOVICH: Yeah.

MR. LEVINE: I think these identity proposals will make forgery more difficult, and it will make the forensics easier to try to determine the actual party responsible for sending an illegal piece of mail, but from the point of view of a Do Not E-mail List, if a Do Not E-mail List lists recipients, which by definition can't be forged rather than senders, I mean, in any Do Not E-mail model that I can think of, it's incumbent upon the sender to obey the Do Not E-mail rules regardless of who he claims to be.

So it's really tangential to the Do Not E-mail issue and even to the whole spam issue.

MR. SHAFRANOVICH: May I take this time to figure out whether these proposals will make tracing of spam easier and prosecution of spam easier?

MS. ROBBINS: I'm sorry, can you repeat that?

MR. SHAFRANOVICH: Are you trying to figure out whether this proposal will make enforcement easier?

MS. ROBBINS: One of the concerns or one of the issues that Congress has asked us to look at is the enforceability of any of these models or of a proposal that we have to propose. We have to talk about the security issue, the privacy issue, the enforceability issues, and that's just one component that we need to look at.

MR. SALSBURG: So if there are authentication systems that are in the works that are being tried out in the marketplace that would assist in enforcement, that bears upon how we're going to evaluate the proposals.

MR. LEVINE: All of these authentication schemes are designed to make it easier to determine the actual sender of a piece of e-mail, and to that extent, yes, it will make enforcement easier since it will basically remove one possible link from the chain to the recipient back to the perpetrator.

I think that in terms of legal issues, it hasn't been a link that's been particularly difficult to follow for people who are motivated enough to sue. The point here is to make it so it can be done automatically by high speed computers, which is a whole separate issue.

MR. SALSBURG: Would the distributed registry model where you put the tag in your DNS information require any changes in protocol?

MR. LEVINE: No.

MR. SALSBURG: It would require some just general agreement that this is where you put the information?

MR. SHAFRANOVICH: Yeah.

MR. LEVINE: Yeah, and the only thing it would require is it would require the software that the mailers use be upgraded to examine the tag before attempting to send the mail, but that would be within the existing protocol, and particularly you would not -- other than publishing the tag, it would not put any burden on the recipients of the mail. Their mail servers would operate exactly as they do now.

MR. SHAFRANOVICH: So it's essentially you're publishing the tag. If you're publishing the tag in the e-mail server, then there will be no changes. That's when the proposal becoming quite complicated, but if you publish the e-mail, the only change that I would think of would be indirectly any record type.

MR. LEVINE: Right, again -- although but it still seems to me that once the recipient networks has published that record, it has no further effect on the way their mail server accepts mail.

MR. SALSBURG: We're going to take a quick pause here for the court reporter.

(Discussion off the record.)

MR. SALSBURG: We're back on. Do either of you have any closing thoughts that you want to provide on the issue of a Do Not E-mail Registry?

MR. LEVINE: For me it's pretty much reiterating what we said before. Technically a domain based system, particularly a distributed domain based system, is straightforward to implement, and I think we've done enough experiments to know technically it could operate.

However, without more effective enforcement which both involves changes in the laws so that recipients can pursue spammers and more funding so that agencies such as yours can go after larger violators on a larger scale, it doesn't matter because spammers have made it pretty clear that their activity's criminal, and without strong enforcement, they're going to ignore whatever nominal rules you attempt to place on them.

MR. SHAFRANOVICH: Yeah. I would just add enforcement is the key. You need funding. You need multiple parties able to sue and strong rules. That's all, it all ties in to enforcing.

MR. SALSBURG: Are there additional people you think we should talk to that you think would help enlighten us?

MR. LEVINE: Based on what I understand -- the answer is yes, but, I'm pretty sure they're already all on your list.

MR. SALSBURG: Okay. We want to thank you, and we're going to turn this over now to our colleague, Julie Bush, who is one of the staff here at the FTC working on the report the Commission has to provide to Congress regarding a bounty system or reward system for catching spammers.

So thank you again, and please feel free to give us a call if you have further comments.

MS. ROBBINS: Thank you very much.

CERTIFICATION OF REPORTER

MATTER NUMBER: P044405
CASE TITLE: INTERVIEWS IN CAN-SPAM REPORT TO CONGRESS
HEARING DATE: FEBRUARY 26, 2004

I HEREBY CERTIFY that the transcript contained herein is a full and accurate transcript of the tapes transcribed by me on the above cause before the FEDERAL TRADE COMMISSION to the best of my knowledge and belief.

DATED: MARCH 11, 2004

DEBRA L. MAHEUX

CERTIFICATION OF PROOFREADER

I HEREBY CERTIFY that I proofread the transcript for accuracy in spelling, hyphenation, punctuation and format.

DIANE QUADE